

**METHOD AND APPARATUS FOR PREPARING A DOCUMENT
TO BE READ BY A TEXT-TO-SPEECH READER**

Inventor(s):

John Brian Pickering

,

International Business Machines Corporation

IBM Docket No. GB9-2001-0104-US1

IBM Disclosure No. GB8-2001-0091

Express Mail Label No.: EV 346749178 US

**METHOD AND APPARATUS FOR PREPARING A DOCUMENT
TO BE READ BY A TEXT-TO-SPEECH READER**

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of United Kingdom Application number 0215123.1, filed June 28, 2002.

BACKGROUND

Field of the Invention

[0002] This invention relates to a method and apparatus for preparing a document to be read by a text-to-speech reader. In particular the invention relates to classifying the text elements in a document according to voice types of a text-to-speech reader.

Description of the Related Art

[0003] In a number of different areas, such as voice access to the Internet, 'reading' textual information for the blind, and creating audio versions of newspapers, there is a significant problem in ensuring that appropriate attention can be drawn to the sections in a given document and the information they contain. One important attentional cue under such circumstances is a change of voice, for instance from male to female voice. In auditory terms, this has the effect of highlighting that something has changed in the informational content.

[0004] Machine-readable documents are a mixture of both mark-up tags, paragraph markers, page breakers, lists and the text itself. The text may further use tags or punctuation marks to provide fine detailed structure of emphasis, for instance, quotation marks and brackets or changing character weight to bold or italic. Furthermore, VoiceXML tags in a document describe how a spoken version should render the structural and informational content.

[0005] One example of such voice-type switching would be a VoiceXML home page with multiple windows and sections. Each window or section line or section of a dialogue may be explicitly identified as belonging to a specific voice.

[0006] A problem with VoiceXML pages is that the VoiceXML tags need to be inserted into a document by the document designer.

[0007] Previously, methods have highlighted grouping content together to drive voice-type selection on the basis of document structure alone. In this way, tables for example can be read out intelligently. However, such systems do not supplement this structuring with thematic information to complete the groupings or the better to select appropriate voice characteristics for output.

SUMMARY OF THE INVENTION

[0008] According to a first aspect of the present invention there is provided a method for preparing a document to be read by a text-to-speech reader. The method can include: identifying two or more voice types available to the text-to-speech reader; identifying the text elements within the document; grouping similar text elements together; and classifying the text elements according to voice types available to the text-to-speech reader.

[0009] Such a solution allows for the automatic population of a document with voice tags thereby voice enabling the document.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] Embodiments of the invention will now be described, by means of example only, with reference to the accompanying drawings in which:

[0011] Figure 1 is a schematic diagram of a source document; a document processor; a voice type characteristic table; and a speech generation unit used in the present embodiment;

[0012] Figure 2 is a schematic diagram of a source document;

[0013] Figure 3 is an example table of voice type characteristics;

[0014] Figure 4 is a flow diagram of the steps in the document processor;

[0015] Figure 5 is an example table of how the source document is classified; and

[0016] Figure 6 is an example of the source document with inserted voice tags.

DETAILED DESCRIPTION OF THE INVENTION

[0017] Referring to Figure 1 there is shown a schematic diagram of a source document 12; a document processor 14; a voice type characteristic table 16; a voice tagged document 18; and a speech generator 20 used to deliver the final speech output 22. The source document 12 and voice type characteristics table 16 are input into the document processor 14. The document 12 is processed and a voice tagged document 18 is output. The speech generator 20 receives the voice tagged document 18 and performs text-to-speech under the control of the voice tags embedded in the document.

[0018] Referring to Figure 2, the example source document 12 is a personal home page 24 comprising three different types of windows. The first and last windows are adverts 26A and 26B, the second window is a news window 28 and the third window is an email inbox window 30. The adverts 26A and 26B in this example are both for a product called Nuts.

[0019] Referring to Figure 3, the voice type characteristic table 16 comprises a column for the voice type identifier 32 and a column for the voice type characteristics 34. In this example voice type 1 is a neutral, authoritative, formal voice like a news reader's; voice type 2 is an informal voice which is friendlier than voice 1; voice type 3 is an enthusiastic voice suitable for advertisements; voice 4 is a particular voice belonging to a personality, in this case the politician quoted in the news item of the news window.

[0020] Referring to Figure 4, a flow diagram of the steps in the document processor is shown. Step 402 identifies all the text elements within the source document 12. Step 404 groups similar text elements together. Step 406 classifies the grouped text elements against the voice type characteristics 34. Step 408 marks up the classified grouped text elements within the source document 12 with voice type identifiers 32. It is this marked-up source document 18 that is passed on to the speech generator.

[0021] Referring to step 402, the identification of all the text elements is performed by a structural parser (not shown). The structural parser is responsible for establishing which sections of the text belong in separate gross sections. It subdivides the complete text into generic sections: this would be analogous to chapters or sections in a book or in this case the separate windows or frames in the document. Gross structural subdivisions such as the frames are marked with sequenced tags <s1> . . . <sN>. Next, individual paragraphs are marked with sequenced tags <p1> . . . <pN>. Next, individual text elements within the paragraph are marked with sequential tags <t1> . . . <tN>. Individual elements include explicit quotations keyed of the orthographic convention of using quotation marks. Also included is a definition keyed off the typographical convention of italicizing or otherwise changing character properties for a run of more than a single word. Further included may be a list keyed by the appropriate mark-up convention, for instance, <o1> . . . </o1> in HTML with each list item marked with .

[0022] The structural parser creates a hierarchical tree showing the text elements and gross sections. In essence, the structural parser simply collates all of the information available from the existing mark-up tags, document structure and document orthography.

[0023] Referring to step 404, the grouping of similar text items together is performed by a thematic parser (not shown) that identifies which of these sections actually belongs together. In the preferred embodiment the thematic parser initially performs a syntactic parse and secondly uses text-mining techniques to group the text elements. In other embodiments step 404 may be performed by either of syntactic parse or text mining. Based on the results of the text mining and syntactic parses, thematic groupings can be made to show which text elements belong to the same topic. In the example given, the two advert frames 26A and 26B need to be linked as they are for the same product or service. If they were for different products or services the same voice type may be used but could be altered to distinguish the two adverts. Alternatively a different voice could be used.

[0024] The inclusion of some degree of syntactic parsing at least for grouping of themes works less efficiently across broader text ranges such as non-sequential paragraphs than it does in the same paragraph. However, it would provide a useful indication of where two non-sequential text elements are related. Take a possible quotation reported in a news broadcast:

[0025] "Our commitment to the people of this area," the politician announced, "has increased in real terms over the last year".

[0026] The structural parser would have identified (based on the opening and closing quotation marks) two text elements: "Our commitment to the people of this area," and "has increased in real terms over the last year". Clearly, however, the latter is simply a continuation of the former, and the two text elements should be treated as dependent. A syntactic parse links these two text elements to be treated as single text element in the remainder of the embodiment. Similarly text elements within sentences without embedded quotations are linked and treated as one. Sentences within a paragraph are similarly linked and treated as one unit.

[0027] The text mining grouping works more efficiently across broader text ranges and, in this embodiment, groups the text elements according to themes found within the text elements. In another embodiment the themes could be a predefined group list such as: adverts, emails, news, and personal. Clearly the pre-defined group list is unlimited. Furthermore, text mining grouping works best with larger sets of words so is best performed after the structural parse.

[0028] The result of the thematic parse is to identify sections of text that belong together, whether they are adjacent or distributed across a document. Each text element from the hierarchical tree is now in a group of similar text elements as shown in Figure 5.

[0029] The set of text elements is input into a clustering program. Altering the composition of the input set of text elements will almost certainly alter the nature and content of the clusters. The clustering program groups the documents in clusters according to the topics that the document covers. The clusters are characterised by a set of words, which can be in the form of several word-pairs. In general, at least one of the word-pairs is present in each document comprising the cluster. These sets of words constitute a primary level of grouping.

[0030] In the described embodiment, the clustering program used is IBM Intelligent Miner for Text provided by International Business Machines Corporation. This is a text-mining tool that takes a collection of text elements in a document and organizes them into a tree-based structure, or taxonomy, based on a similarity between meanings of text elements.

[0031] The starting point for the IBM Intelligent Miner for Text program are clusters which include only one text element and these are referred to as "singletons". The program then tries to merge singletons into larger clusters, then to merge those clusters into even larger clusters, and so on. The ideal outcome when clustering is complete is to have as few remaining singletons as possible.

[0032] If a tree-based structure is considered, each branch of the tree can be thought of as a cluster. At the top of the tree is the biggest cluster, containing all the text-elements. This is subdivided into smaller clusters, and these into still smaller clusters, until the smallest branches that contain only one text element (or effective text element). Typically, the clusters at a given level do not overlap, so that each text element appears only once, under only one branch.

[0033] The concept of similarity of text elements requires a similarity measure. A simple method would be to consider the frequency of single words, and to base similarity on the closeness of this profile between documents. However, this would be noisy and imprecise due to lexical ambiguity and synonyms. The method used in IBM's Intelligent Miner for Text program is to find lexical affinities within the text element. In other words, correlations of pairs of words appearing frequently within short distances throughout the document.

[0034] A similarity measure is then based on these lexical affinities. Identified pairs of terms for a text element are collected in term sets, these sets are compared to each other and the term set of a cluster is a merge of the term sets of its sub-clusters.

[0035] Other forms of extraction of keywords can be used in place of IBM's Intelligent Miner for Text program. The aim is to obtain a plurality of sets of words that characterise the concepts represented by the text elements.

[0036] Referring to step 406, the classifying of the grouped text elements against voice types is performed by a pragmatic parser (not shown). The pragmatic parser matches each group of text elements to a voice type characterisation using a text comparison method. In the preferred embodiment this method is Latent Semantic Analysis (LSA) again performed by IBM Intelligent Miner for Text. With LSA each existing group of text elements is classified using the voice types as categories. Having keywords in the voice type characterisation 34 helps this process.

[0037] In the preferred embodiment keywords for the type of text element grouping are used. For instance, putting the words "news reader, news item, news article" in the voice type classification 34 for voice type 1 helps the classifying process match news articles against voice type 1 which is suitable for reading news articles. Other types would include adverts, email, personal column, reviews, and schedules. These keywords are placed in the voice type characterisation 34 for the particular voice that the words refer to.

[0038] In another embodiment the pragmatic parser will look for intention in the text element groups and intentional words are placed in the voice type characterisation 34. For instance, voice one is characterised as neutral, authoritative and formal, the LSA will match the text element grouping that best fits this characterisation.

[0039] Voice type 5 is a special case of the type of text element grouping. Voice type 5 impersonates a particular politician and the politician's name is in the voice type characterisation 34. The thematic parser will pick up if a particular person says the quotations and the pragmatic parser will match the voice to the quotation.

[0040] Latent Semantic Analysis (LSA) is a fully automatic mathematical/statistical technique for extracting relations of expected contextual usage of words in passages of text. This process is used in the preferred embodiment. Other forms of Latent Semantic Indexing or automatic word meaning comparisons could be used.

[0041] LSA used in the pragmatic parser has two inputs. The first input is a group of text elements. The second input is the voice type characterisations. The pragmatic parser has an output that provides an indication of the correlation between the groups of text elements and the voice type characterisations.

[0042] Although a reader does not need to understand the internal process of LSA in order to put the invention into practice, for the sake of completeness a brief overview of the LSA process within the automated system is given.

[0043] The text elements of the document form the columns of a matrix. Each cell in the matrix contains the frequency with which a word of its row appears in the text element. The cell entries are subjected to a preliminary transformation in which each cell frequency is weighted by a function that expresses both the word's importance in the particular passage and the degree to which the word type carries information in the domain of discourse in general.

[0044] The LSA applies singular value decomposition (SVD) to the matrix. This is a general form of factor analysis that condenses the very large matrix of word-by-context data into a much smaller (but still typically 100-500) dimensional representation. In SVD, a rectangular matrix is decomposed into the product of three other matrices. One component matrix describes the original row entities as vectors of derived orthogonal factor values, another describes the original column entities in the same way, and the third is a diagonal matrix containing scaling values such that when the three components are matrix-multiplied, the original matrix is reconstructed. Any matrix can be so decomposed perfectly, using no more factors than the smallest dimension of the original matrix.

[0045] Each word has a vector based on the values of the row in the matrix reduced by SVD for that word. Two words can be compared by measuring the cosine of the angle between the vectors of the two words in a pre-constructed multidimensional semantic space. Similarly, two text elements each containing a plurality of words can be compared. Each text element has a vector produced by summing the vectors of the individual words in the passage.

[0046] In this case the text elements are a set of words from the source document. The similarity between resulting vectors for text elements, as measured by the cosine of their contained angle, has been shown to closely mimic human judgments of meaning similarity. The measurement of the cosine of the contained angle provides a value for each comparison of a text element with a source text.

[0047] In the pragmatic parser a set of voice type characterisation words and a group of text elements are input into an LSA program. For example, the set of words "neutral, authoritative, formal" and the words of a particular text element group are input. The program outputs a value of correlation between the set of words and the text element group. This is repeated for each set of voice characterisations and for each text element group text in a one to one mapping until a set of values is obtained.

[0048] Referring to Figure 5, the grouping of the text elements after processing is shown followed by the classification. The first grouping is the news narrative in the Local News Window 28 which is classified with voice type 1. The next grouping is the statements by the politician classified by voice type 4. The next grouping is the statement made by the opposition for which there is no set voice and voice type 1* is used. In this case the nearest voice is matched and marked with a "*" to indicate that a modification to the voice output should be made when reading to distinguish it from nearest voice.

[0049] Modification would be effected as follows. For a full TTS system for speech output, the prosodic parameters relating to segmental and supra-segmental duration, pitch and intensity would be varied. If the mean pitch is varied beyond half an octave then distortion may occur so normalization of the voice signal would be effected. For pre-recorded audio output, the source characteristics of, for instance, Linear Predictive Coding (LPC) analysis would be modified in respect of pitch only, limited to mean pitch value differences of a third an octave.

[0050] The next grouping is the text in the Email Inbox Window 30 and voice type 2 is assigned. The last grouping is the adverts 26A, 26B and voice type 3 is assigned to both adverts which are treated as one text element.

[0051] Referring to Figure 6, the voice tags are show between '<' '>' symbols. The adverts both have <voice3> tags preceding them. The email window has a <voice2> tag preceding the text. The Local News window has a mixture of <voice1>, <voice1*> and <voice4> tags.